

IN THE UNITED STATES PATENT AND TRADEMARK OFFICE

APPLICATION FOR LETTERS PATENT

EV205822750

**Midstream Determination of Varying
Bandwidth Availability**

Inventors:

Troy D. Batterberry
Alexandre Grigorovitch
Anders E. Klemets
James C. Stewart
David J. Roth
Ye Jin Choi

ATTORNEY'S DOCKET NO. MS1-1479US

TECHNICAL FIELD

The systems and methods described herein relate to measuring bandwidth availability. More particularly, the systems and methods described herein relate to midstream determination of bandwidth availability for a connection between to entities on a network, even in cases where available bandwidth varies.

BACKGROUND

As the Internet has matured, the format characteristics of the content available on the Internet have changed. Sound and video content is now mixed in with the traditional textual content. However, this new content on the Internet requires a greater connection speed (i.e., bandwidth) than was commonly available a few years ago.

Fig. 1 illustrates an example of a typical Internet configuration. It includes a server (such as media server 20), which is coupled to the Internet 30. The server typically includes one or more physical server computers 22 with one or more physical storage devices and/or databases 24. On the other side of an Internet transmission is a client 90, which is connected via one of many available Internet Service Providers (ISPs) 80. Herein, a server is a network entity that sends data and a client is a network entity that receives data.

Cloud 30 is labeled the Internet, but it is understood that this cloud represents that portion of the Internet that does not include the server, client's ISP, and the client. Inside such cloud are the routers, transmission lines, connections, and other devices that more-often-than-not successfully transmit data between clients and servers. Inside exemplary Internet cloud 30 are routers 32-44; two

1 satellite dishes 46 and 50; and a satellite 48. These represent the possible paths
2 that a data packet may take on its way between the server and the client.

3 **Bandwidth**

4 Bandwidth is the amount of data that can be transmitted in a fixed amount
5 of time. For example, bandwidth between media server 20 in Fig. 1 to media
6 client 90 is calculated by the amount of data (e.g., 1000 bits) that may be
7 transmitted between them in a unit of time (e.g., one second).

8 As shown in Fig. 1, a transmission over the Internet travels across multiple
9 links before it reaches its destination. Each link has its own bandwidth. Like a
10 chain being only as strong as its weakest link, the maximum bandwidth between
11 server 20 and client 90 is the link therebetween with the slowest bandwidth.
12 Typically, that is the link (such as link 82 in Fig. 1) between the client 90 and its
13 ISPs 80. That slowest bandwidth is the maximum de facto bandwidth.

14 Herein, unless otherwise apparent from the context, references to
15 bandwidth between network entities (such as server 20 and client 90) is assumed
16 to be the maximum de facto bandwidth therebetween.

17 Bandwidth may also be called “connection speed”, “speed”, or “rate”. In
18 references to bandwidth measured by bits per second, it may also be called “bit
19 rate” or “bitrate.”

20 **Streaming Media**

21 Streaming is a technique for transferring multimedia data such that it can be
22 processed as a steady and continuous stream. Streaming technologies are
23 becoming increasingly important with the growth of the Internet because most
24 users do not have fast enough access to download large multimedia files quickly.
25

1 With streaming, the client browser or plug-in can start displaying the data before
2 the entire file has been transmitted.

3 For streaming to work, the client side receiving the data must be able to
4 collect the data and send it as a steady stream to the application that is processing
5 the data and converting it to sound or pictures. This means that if the streaming
6 client receives the data more quickly than required, it needs to save the excess data
7 in a buffer. If the data doesn't come quickly enough, however, the presentation of
8 the data will not be smooth.

9 Within the context of an audio and/or visual presentation, "media" and
10 "multimedia" are used interchangeably herein. Media refers to the presentation of
11 text, graphics, video, animation, and/or sound in an integrated way.

12 "Streaming media" is an audio and/or visual presentation that is transmitted
13 over a network (such as the Internet) to an end-user. Such transmission is
14 performed so that the presentation is relatively smooth and not jerky. Long pauses
15 while additional frames are being downloaded to the user are annoying to the user.
16 These annoyances encourage a user to avoid viewing future streaming media.

17 **Smoothly Transmitting Streaming Media**

18 Since the bandwidth determines the rate at which the client will receive
19 data, a streaming media presentation may only be presented at a rate no greater
20 than what the bandwidth allows. For example, assume media server 20 needs to
21 send data at 50Kbps to the client 90 in order to smoothly "play" a streaming media
22 presentation. However, the bandwidth between the client and server is only
23 30Kbps. The result is a jerky and jumpy media presentation.

24 In an effort to alleviate this problem, streaming media presentations are
25 often encoded into multiple formats with differing degrees of qualities. The

1 formats with the lowest quality (e.g., small size, low resolution, small color
2 palette) have the least amount of data to push to the client over a given time.
3 Therefore, a client over a slow link can smoothly present the streaming media
4 presentation, but the quality of the presentation suffers. The formats with the
5 highest quality (e.g., full screen size, high resolution, large color palette) have the
6 greatest amount of data to push to the client over a given time. Therefore, the
7 client with a fast link can smoothly present the streaming media presentation and
8 still provide a high quality presentation.

9 **Select-a-Bandwidth Approach**

10 When a server sends streaming media to a client, it needs to know what
11 format to use. Thus, in order to select the proper format, the server must to know
12 the bandwidth between the server and the client.

13 This easiest way to accomplish this is to ask the user of the client what their
14 bandwidth is. Since a client's link to the Internet is typically the bandwidth
15 bottleneck, knowing the bandwidth of this link typically indicates the actual
16 bandwidth.

17 Fig. 2 shows a cut-away 100 of a Web page displayed on a client's
18 computer. Inside the cut-away 100, is a typical user-interface 110 that may be
19 used to ask a user what their connection speed is. The user clicks on one of the
20 three buttons 112, 114, and 116 provided by the user-interface 110. If the user
21 clicks on button 112, the server delivers data from a file containing streaming
22 media in a format designed for transmission at 28.8 Kbps. Likewise, if the user
23 clicks on button 114, data sends from a file containing streaming media in a format
24 designed for transmission at 56.6 Kbps. If the user clicks on button 114, the server
25 delivers data from a file containing streaming media in a format designed for

1 transmission at a rate greater than 56.6Kbps and up-to the typical speed of a T1
2 connection.

3 However, the primary problem with the "select-a-bandwidth" approach is
4 that it requires a thoughtful selection by a user. This approach invites selection
5 errors.

6 It requires that a user care, understand, and have knowledge of her
7 connection speed. Often, a user does not pay particular attention to which button
8 to press. The user may only know that a media presentation will appear if the user
9 presses one of these buttons. Therefore, they press any one of them.

10 Often, a user does not understand the concept of bandwidth. A user may
11 choose button 116 because she may want to see the presentation at its highest
12 quality. This user does not realize that seeing the presentation at its highest quality
13 may result in a non-smooth presentation because her Internet connection cannot
14 handle the rate that the data is being sent through it.

15 If she does understand the concept of bandwidth, then the user may not
16 know her bandwidth. A user may simply be ignorant of her bandwidth. In addition,
17 varying degrees of noise may cause varying connection speeds each time a user
18 connects to the Internet. Furthermore, some types of connections (such as a cable
19 modem) can have wide degrees of connection speed depending upon numerous
20 factors.

21 Moreover, the user needs to understand the implications of an incorrect
22 choice. A user needs to be educated so that she understands that she needs to select
23 an option that is equal to or less than her bandwidth to get a smooth presentation.
24 But she should not choose one that is significantly less than her bandwidth. If she
25

1 does, then she will be seeing a smooth presentation at a lower quality that she
2 could otherwise see at a higher available bandwidth.

3 As can be seen by the above discussion, this manual approach is often
4 confusing and intimidating to many user. Therefore, it often results in incorrect
5 selections.

6 What's more, maintaining multiple files (one for each bandwidth) at the
7 media server adds to the overhead of maintaining a Web site.

8 **Automatic Bandwidth Detection**

9 To overcome these problems, media servers use a single file containing
10 subfiles for multiple bandwidths. Also, media servers automatically detect the
11 bandwidth.

12 This single file is called a MBR (multiple bit rate) file. The MBR files
13 typically include multiple differing "bands" or "streams." These bands may be
14 called "subfiles." A user only clicks on one link. Automatically, behind the scenes,
15 the server determines the proper stream to send to the client based on the speed
16 selected by the client.

17 In an environment where end-to-end latency is very high, this automatic
18 speed detection may take a long time. This means that an additional five to thirty
19 seconds is added to the user's wait for the presentation to begin. One factor in this
20 delay for existing automatic speed detection is because of long "handshaking"
21 times while the speed determination is going on.

22 One existing automatic detection technique involves sending multiple data
23 packets for measuring the speed between the server and client. This technique is
24 described further below in the section titled, "Multiple Measurement Packets
25 Technique."

Bandwidth Measurement Packets

Typically, automatic bandwidth detection techniques measure bandwidth between entities on a network by sending one or more packets of a known size.

Fig. 3 shows a time graph tracking the transmission of two such packets (P_x and P_y) between a sender (e.g., server) and a receiver (e.g., client). The server and client sides are labeled so. On the graph, time advanced downwardly.

Time t_a indicates the time at the server the transmission of P_x begins. Time t_b indicates the time at the server the transmission of P_x ends. Similarly, Time t_0 indicates the time at the client begins receiving P_x . Time t_1 indicates the time at the client completes reception of P_x . At t_1 , the network hardware presumably passes the packet up the communication layers to the application layer.

Packet P_y is similarly labeled on the time graph of Fig. 3. t_c is the server time at the transmission of P_y begins. t_d is the server time that the transmission of P_y ends. Similarly, t_2 the client time that it begins receiving P_y . t_3 is the client time that it completes reception of P_y . At t_3 , the network hardware presumably passes the packet up the communication layers to the application layer.

Bandwidth measurement using a single packet.

In a controlled, laboratory-like environment, measuring bandwidth between two entities on a network is straightforward. To make such a calculation, send a packet of a known size from one entity to the other and measure the transmission latency, which is the amount of time it takes a packet to travel from source to destination. Given this scenario, one must know the time that the packet was sent and the time that the packet arrived.

This technique is nearly completely impractical outside of the laboratory setting. It cannot be used in an asynchronous network (like the Internet) because it

1 requires synchronization between the client and server. Both must be using the
2 same clock.

3 Alternatively, the client may track the time it begins receiving a packet
4 (such as t_0 for P_x) and the time the packet is completely received (such as t_1 for
5 P_x).

6 Fig. 3 shows packet P_x being sent from a server to a client. P_x has a known
7 size in bits of PS. The formula for calculating bandwidth (bw) is

$$\text{bw}(P_x) = \frac{\text{PS}}{t_1 - t_0}$$

8
9
10
11 **Formula 1 (Single Packet)**

12 This technique works in theory, but unfortunately does not work in practice.
13 Only the hardware knows when a packet is initially received. Therefore, only the
14 hardware knows when t_0 is.

15 The other communication layers (such as the transport layer and the
16 application layer) can only discover the time when the packet is completely
17 received by the hardware. That is when the hardware passes it up to them. This
18 completion time for packet P_x is t_1 . It is not possible to calculate bandwidth only
19 one knowing one point in time.

20 Packet-pair. A technique called packet-pair is used to overcome these
21 problems in asynchronous networks. With packet-pair, the server sends a pair of
22 packets, one immediately after the other. The bandwidth is determined by
23 dividing the packet size by the time difference in reception of each packet.

24 Each packet has specific measurable characteristics. In particular, these
25 characteristics include its packet size (PS) and the measured time such a packet

1 arrives (e.g., $t_{0.3}$ in Fig. 3). Some characteristics (such as packet size) may be
2 specified rather than measured, but they may be measured if so desired.

3 As shown in Fig. 3, the server sends packet, P_x . The client's hardware
4 begins receiving the packet at t_0 . When reception of the packet is complete at t_1 ,
5 the hardware passes it up the communication layers. Ultimately, it is received by
6 the destination layer (e.g., application layer) at presumably t_1 .

7 After the server sends P_x (which is completed at t_b), it immediately sends
8 packet P_y at t_c . It is important that there be either 1) absolutely no measurable
9 delay between t_b and t_c or 2) a delay of a known length between t_b and t_c . Herein,
10 to simplify the description, it will be assumed that there is no measurable delay
11 between t_b and t_c .

12 The client's hardware begins receiving P_y at t_2 . When reception of the
13 packet is complete at t_3 , the hardware passes it up the communication layers.
14 Ultimately, it is received by the destination layer (e.g., application layer) at
15 presumably t_3 .

16 Fig. 3 shows no delay between t_1 (the time of completion of reception of P_x)
17 and t_2 (the time reception of P_y begins). Theoretically, this will always be the case
18 if P_x and P_y are transmitted under identical conditions. In practice, is the often the
19 case because P_y is sent immediately after P_x .

20 Using packet-pair, the formula for calculating bandwidth (bw) is

$$21 \quad bw(P_x P_y) = \frac{PS}{t_3 - t_1}$$

22
23 **Formula 2 (Packet-Pair)**
24
25

1 This technique works in theory and in practice. However, it only works
2 well over a network that is relatively static.

3 For example, in Fig. 1, assume the network consists of only the server 20;
4 routers 32, 34, and 36; a specific ISP of ISPs 80; and client 90. Further, assume
5 that the links between each node on this static network is fixed and has a
6 consistent bandwidth. In this situation, the packet-pair techniques provide an
7 accurate and effective measurement of bandwidth.

8 Packet-pair does not work well over the Internet. However, the packet-pair
9 technique does not work well over a dynamic network, like the Internet. A
10 dynamic network is one where there is a possibility that a packet may be handled
11 in a manner different from an earlier packet or different from a later packet.

12 Fig. 1 illustrates examples of those handling differences. Assume that all
13 packets are traveling from the server to the client (from left to right in Fig. 1).
14 Assume that packets 60-68 were sent back-to-back by the server 20 to the client
15 90. Assume that packet 70 was sent by another server (not shown) to the client 90
16 and it is unrelated to bandwidth measurement.

17 Notice, as illustrated in Fig. 1, that packets may take different routes. In
18 addition, some routes may significantly delay the packet transmission. This is
19 especially true if the packet is transmitted via an apparently unusual (but not
20 necessarily uncommon) route, such as wireless transmission, overseas via an
21 underwater cable, satellite transmission (as shown by dishes 46 and 50 and
22 satellite 48), etc.

23 A router (such as router 42) may delay a packet (such as 64) more than
24 another may by temporarily buffering it. Another packet (such as packet 70) from
25

1 another source may slip in between two packets (such as packets 60 and 62). In
2 addition, a modem (not shown) of the client may compress packets.

3 Communications equipment (such as a modem) may compress a packet
4 (such as 66) to shrink the packet size and thus speed along transmission. Such
5 packet compression can significantly affect the bandwidth measurement because
6 not all of the subsequent data packets will be compressed or compressed at the
7 same rate.

8 **Multiple Measurement Packets Technique**

9 To overcome these problems, conventional automatic bandwidth
10 measurement techniques uses multiple packets. A server sends several (much more
11 than two) packets and calculates the speed of each. Conventional wisdom on
12 bandwidth measurement indicates that in order to get accurate measurements
13 several pairs of packets must be sent repeatedly over several seconds to several
14 minutes. Herein, this technique is called “multiple-packets” to distinguish it from
15 the above-described “packet-pair” technique.

16 Typically, the ultimate bandwidth is determined by finding the average of
17 the many bandwidth measurements. This averaging smoothes out variances in
18 delays for each packet; however, it does not compensate for packet compression
19 during transmission. One of two extremely incorrect measurements will skew the
20 average.

21 Unfortunately, this technique takes a long time relative the existing wait for
22 the user between click and media presentation. A long time may be five seconds
23 to several minutes depending on the data and the situation. Such a delay adds to
24 the annoyance factor for the user who wishes experience the media presentation.
25

1 This is not an acceptable delay. Since there are no other options available using
2 conventional techniques, the user is forced to endure these delays.

3 Moreover, these conventional approaches typically use TCP to transmit the
4 packets. Using TCP introduces additional delays for handshaking. These
5 conventional approaches typically modify the kernel of the operating system
6 (usually the transport layer) to perform these measurements.

7 **Varying Bandwidth**

8 Another problem encountered with streaming multimedia content is that,
9 after the initial bandwidth measurement is taken and a streaming rate is
10 determined, factors that influenced the selection of the streaming rate may change.
11 As a result, the bandwidth that is available for streaming media changes as well.
12 For example, if network congestion was a problem at the time the measurement
13 was taken, then the selected streaming rate may be lower than it could be when the
14 network congestion clears. Conversely, if network congestion occurs after the
15 streaming rate has been selected, then the quality of the multimedia presentation
16 may suffer as a result of streaming at a faster rate than the network can
17 accommodate.

18 Sole reliance on initial bandwidth measurements, therefore, may provide a
19 less than optimum streaming experience.

21 **SUMMARY**

22 Various systems and methods described herein provide for midstream
23 determination of varying available bandwidth for streaming content between two
24 network entities.
25

1 In at least some implementations, during content streaming, a client may
2 request a server to send data at an increased transmission rate for a limited period
3 of time, i.e. to “surge” the content transmission. One or more bandwidth
4 measurements are taken during the surge to determine if the increased
5 transmission rate is more optimally suited for the content stream.

6 If the transmission surge is determined to be viable for transmission of the
7 content stream, the client may request the server to continue to transmit remaining
8 content at the higher transmission rate. If the higher transmission rate is not
9 suitable for future transmission, then the rate may remain at the original
10 transmission rate.

BRIEF DESCRIPTION OF THE DRAWINGS

The same numbers are used throughout the drawings to reference like features and components.

Fig. 1 illustrates a typical public networking environment (such as the Internet) and the routing of and delay of data packets sent from a server to a client.

Fig. 2 is cut-away portion of a Web page. The cut-away shows a user interface providing a user a mechanism for selecting the bandwidth. This shows a conventional technique for determining bandwidth.

Fig. 3 shows a packet pair (being sent from a server to a client) graphed in the time domain. This shows a conventional implementation of packet-pair technique to measure bandwidth.

Fig. 4 is a visual representation of a history list of the last ten measured bandwidths.

Fig. 5 is a block diagram of an exemplary server-client environment for midstream determination of available streaming bandwidth.

Fig. 6 is a flowchart illustrating the methodology of an implementation of the server side of the exemplary midstream bandwidth measurement technique.

Fig. 7 is a flowchart illustrating the methodology of an implementation of the client side of the exemplary midstream bandwidth measurement technique.

Fig. 8 is an example of a computing operating environment capable of implementing the server side and/or the of the exemplary bandwidth meter.

DETAILED DESCRIPTION

Systems and methods for midstream determination of varying network bandwidth for streaming content are described herein. One or more exemplary implementations of systems and methods for determining bandwidth availability during content streaming are shown. The described systems and methods are exemplary only and are not meant to limit the scope of the appended claims. Other embodiments not described herein may be implemented without departing from the scope of the appended claims.

Utilizing Content Packets for Packet Pair Measurements

Using back to back data packets to measure throughput and bandwidth availability is described in U.S. Patent Application Number 09/636,456, filed 08/09/2000, entitled "Fast Dynamic Measurement of Connection Bandwidth" and assigned to MICROSOFT CORP®, the assignee of the present application. Said patent application is hereby incorporated by reference. Similar logic is extended in the presently described systems and methods to include sending back to back content packets (where "back to back" means without any significant intervening time gap) while actually streaming to determine instantaneous bandwidth and therefore allow the client to more quickly act upon link congestion.

The described systems and methods also use this technique to determine availability of additional bandwidth. In addition to using occasional single packet pairs during the streaming process, multiple packet pairs can be sent back to back to develop both individual packet pair measurements as well as an aggregate measurement of the total data received over the span of the packets.

For example, consider a case in which four 1500 byte packets are sent immediately subsequent to each other. In addition to yielding three individual

1 packet pair measurements, the time required to receive all three packets can also
2 be combined to allow a more precise measurement of bandwidth over a larger
3 timeframe.

4 Since all of these measurements take place using content that is in the
5 process of being streamed anyway, no additional bandwidth needs to be consumed
6 and no additional latency is incurred. Since most encoded content is already
7 highly entropic, compression algorithms used in the streaming process will not
8 skew the measurement results. Only the server needs to flag the packets as being
9 back to back so that the client recognizes that it can use them for dynamic
10 (midstream) bandwidth measurements.

11 **History Lists**

12 Fig. 4 is a visual representation of a history list 300 of recorded bandwidth
13 measurements. A client keeps this history of calculated bandwidths to ameliorate
14 the effect of an inaccurate bandwidth measurement. Upon re-connection with a
15 server to which the client has previously been connected, a client will access the
16 history list 300 and initially connect at a rate equal to the median rate included in
17 the history list 300. It is noted that the client could determine a connection rate
18 from the history list 300 in another manner, such as by taking a mean of the rates
19 included in the history list 300. However, utilizing a median rate from the history
20 list 300 is less likely to be skewed by an extremely aberrant entry.

21 In at least one other implementation, the history list 300 may be accessed
22 upon connection with any server, not just a server to which the client has
23 previously been connected.

24 Items are added and removed from the list in a FIFO (first in, first out)
25 method. Utilizing the FIFO technique, the latest ten measurements are always

1 included in the list (although the history list 300 may contain a greater or fewer
2 number of rates than ten).

3 In the systems and methods described herein, bandwidth measurements
4 taken mid-stream are stored in a history list. Techniques that take bandwidth
5 measurements at an initial connection between client and server may not be
6 entirely accurate.

7 For example, consider a user of a proxy device that marshals all TCP traffic
8 through a latent architecture prior to forwarding on to the client. While the
9 firewall may be capable of processing relatively high bitrate streams, it may
10 inadvertently add time between the forwarding of a specific packet and therefore
11 cause the client to believe it has an artificially low (or high) bandwidth connection.
12 If the history lists only contain the initial measurement and if the behavior is
13 consistent between attempts (as is the case with many proxies/firewalls), the client
14 will always believe it has an artificially low (or high) bandwidth connection.

15 By taking midstream measurements of available bandwidth, the client can
16 correct the values stored in the history lists. By using a history list, a good
17 estimation of the link bandwidth may be determined very quickly, thus providing a
18 high-quality experience for a user. Further functionality of the history list 300 will
19 be described in greater detail, below.

20 **Exemplary Environment**

21 Fig. 5 is a block diagram of an exemplary server-client environment 600 for
22 midstream determination of available streaming bandwidth. The environment 600
23 includes a server 602 and a client 604 that communicate with each other through a
24 network 606. The network 606 may be the Internet, a local area network, a wide
25 access network, or the like.

1 The server 600 includes a processor 608, a network interface 610 for
2 communicating with the network 606, and memory 612. The memory stores a
3 server operating system 614, a control module 616 and one or more multi-bitrate
4 (MBR) files 618. In the present example, the MBR files 618 contain multimedia
5 data that may be streamed to the client 604. The control module 616 oversees
6 execution of the server functionality described below to determine bandwidth
7 availability during the streaming process.

8 The client 604 includes a processor 620, a network interface 622, a user
9 input module 624 (e.g. a keyboard, mouse, etc.) and a display 626. The client 604
10 also includes memory 630, which stores a client operating system 632, a
11 bandwidth measurement module 634, a history module 636 and one or more
12 applications 638 that utilize streaming content from one or more remote source.

13 The control module 634 controls the execution of client functionality for
14 measuring available bandwidth during content streaming and will be discussed in
15 greater detail below. The bandwidth measurement module 636 provides the
16 functionality to calculate the bandwidth based on the packet size, time sent, time
17 received, etc.

18 The history module 638 stores a history list (similar to the history list 300
19 shown in Fig. 4) for transmission rates experienced with each of several sources
20 from which the client 604 has previously received streaming content, the history
21 lists being designated herein as source_1 642, source_2 644 through source_n 646.

22 The memory 630 also stores a content buffer 648 that is used to store
23 content received from the server 602 until the client 604 is ready to play the
24 content. Content buffers are well known in the art and are used to prevent artifacts
25 in the content presentation.

1 The elements and/or modules shown and described with respect to Fig. 5
2 may be implemented as hardware, software or a combination of both. The
3 functionality of the elements and/or modules as pertaining to the appended claims
4 will be described in greater detail, below, with respect to subsequent figures.

5 **Methodological Implementation – Server**

6 Fig. 6 is a flowchart illustrating the methodology of an implementation of
7 the server side of the exemplary midstream bandwidth measurement techniques
8 described herein. In the discussion of Fig. 6, continuing reference will be made to
9 the elements and reference numerals of Fig. 5. Unless otherwise noted, the
10 functionality attributed to the server 602 below is carried out by the control
11 module 616 of the server 602.

12 At block 700, the client 604 connects to the server 602 at an initial
13 transmission rate (Rate R1) to receive content streaming from an MBR file 618.
14 Rate R1 may be determined by any method known in the art, whether or not
15 described above in the Background section. The initial transmission rate may be
16 user-selected or may be automatically determined upon connection of the client
17 604 to the server 602.

18 At some point during the transmission of the streaming content, the server
19 602 receives a request from the client 604 to surge the transmission rate to a rate
20 higher than Rate R1 (block 702). The duration of the surge may be specified by
21 the client in one of several ways.

22 For example, the request may include a command to send a particular
23 number of seconds of data (content) at a higher rate, to send a certain number of
24 data packets at a higher rate, or to send a specified number of bytes of data at a
25 higher rate. Any method known in the art by which the client 604 may request a

1 particular amount of surged data from the server 602 may be implemented without
2 departing from the scope of the described systems and methods.

3 As long as no surge request is detected by the server 602 (“No” branch,
4 block 702) then the streaming continues to the client 604 at Rate R1 (block 700).
5 When a surge request is received by the server 602 (“Yes” branch, block 702),
6 then the server 604 transmits the appropriate amount of data to the client 604 at a
7 higher transmission rate – Rate R2 - at block 704. The higher transmission rate,
8 Rate R2, may be specified by the client or may be previously set by agreement
9 between the client and server entities, such as an agreement that a surge request
10 denotes a request to surge data at a rate that is 10% higher than the current
11 transmission rate or to increase the rate to a next higher rate in a series of
12 predetermined rates.

13 Rate R2 need not be the next higher streaming bandwidth in the streamed
14 MBR file 618; the transmission rate may be increased gradually to determine the
15 viability of switching to the next higher streaming bandwidth. For example, if the
16 client 604 desires to switch from 100 Kbps to 300 Kbps, the client 604 may
17 initially request a surge rate of 200 Kbps. By gradually increasing the
18 transmission bitrate, the client can observe when congestion begins to occur
19 without necessarily flooding the link entirely and causing a bad end user
20 experience.

21 During a surge, the server 602 does not necessarily switch from
22 transmitting a first stream in an MBR file 618 to a second stream in the MBR file
23 618. This is because typical streaming clients can only start rendering on discrete
24 points such as frame boundaries or key frames. When a client shifts to another
25 stream in a conventional MBR file, it must shift on a key-frame boundary for the

1 end user experience to be seamless. Key-frame boundaries are usually not aligned
2 between different streams in an MBR file. Since key-frames occur as little as once
3 every 8-20 seconds depending on the encoded bitrate of the file, a client that
4 incorrectly shifts to a higher bandwidth stream may have a noticeable adverse
5 impact to the viewing experience, since it must wait for another key frame to
6 occur on the destination stream before it can begin re-buffering again when
7 switching back down.

8 To mitigate this problem, it is valuable to first verify that the network
9 connection can sustain the increased bandwidth before actually switching to a
10 different stream. The client 604 can do this by requesting the existing stream to be
11 sent at a rate equivalent to the higher bitrate stream. If the bandwidth is
12 insufficient, the client can remain on the same (lower) bitrate stream and therefore
13 avoid waiting for the arrival of the next key frame had it erroneously switched up
14 and then back down again.

15 After the surged data has been sent by the server 602 to the client 604, the
16 server 602 resumes transmission at Rate R1 (block 706). Thereafter, at block 708,
17 the server 602 may receive a request from the client 604 to begin streaming data
18 from a different stream in the multiple bit rate file ("Yes" branch, block 708), i.e.
19 to increase (or decrease) the current transmission rate (Rate R1). The server 602
20 then changes streams at block 710 and computes a new transmission rate, Rate R1,
21 that corresponds to the newly selected stream. The new Rate R1 is not necessarily
22 the same as the surged transmission rate, Rate R2.

23 As long as a stream change request is not received ("No" branch, block
24 708), then the server 602 continues to stream at Rate R1 (block 700) while
25

1 monitoring for another surge request (block 702) or stream change request (block
2 708).

3 It is noted that there is not necessarily a one-to-one correspondence
4 between surge requests and stream change requests. Also, the order of a surge
5 request and a stream change request is not necessarily as shown in the flow
6 diagram of Fig. 6, though the order shown is that which will typically be
7 encountered in practice. For example, it is possible for a client to send multiple
8 surge requests before it sends a stream change request. It is also possible for a
9 client to send a stream change request to a lower-rate stream (i.e. a request to
10 decrease the transmission rate) immediately after sending a stream change request
11 to a higher-rate stream without sending an intermediate surge request.

12 It is noted that the previous example deals with surging the transmission
13 rate to take advantage of abatement in network congestion. Those skilled in the art
14 will readily understand that the system may also be used to request a decrease in
15 the transmission rate in response to network congestion.

16 **Methodological Implementation – Client**

17 Fig. 7 is a flowchart illustrating the methodology of an implementation of
18 the client side of the exemplary midstream bandwidth measurement technique. In
19 the discussion of Fig. 7, continuing reference will be made to the elements and
20 reference numerals of Fig. 5.

21 At block 800, the user input module 624 of the client 604 receives a
22 command to connect to the server 602 to receive streaming content. This input
23 may simply be a mouse click on a representation of an MBR file 618 on the server
24 602. The control module 634 determines if the client 604 has previously
25 connected to the server 602 at block 802. If not (“No” branch, block 802), then

1 the control module 634 determines the initial rate at which to connect to the server
2 at block 804. This may be done in any way described herein or known in the art.

3 If the client 604 has previously connected to the server 602 ("Yes" branch,
4 block 802), then the control module 634 refers to the history module 638 to
5 determine an appropriate rate at which to initially connect to the server 602. More
6 particularly, a history file associated with the server (source_1 642 ... source_n
7 646) is examined and the initial rate is set at the median rate of the transmission
8 rates included in the history file associated with the server 602 (at block 806).
9 Taking the median of the stored measurements is not necessary, as other means of
10 determining the initial rate may be implemented, such as by taking the mean of the
11 rates. However, for this example, the median rate is used as the initial connection
12 rate.

13 At block 808, the client 604 connects to the server 602 at the initial rate and
14 begins to receive streaming content from the server 602 at block 810. As long as
15 the client 604 does not determine it is time to request a transmission rate surge
16 ("No" branch, block 812), the client continues to receive the streaming data (block
17 810) at the initial, or current, rate as long as there is more streaming data to receive
18 ("Yes" branch, block 824).

19 The determination of when a surge is desired may be based on one or more
20 criteria programmed into the client 604. In at least one implementation, the
21 control module 634 is configured to request a rate surge periodically to determine
22 if the streaming rate can be increased.

23 When the control module 634 determines that it is time to request a surge
24 ("Yes" branch, block 812), the client 604 requests the server 602 to surge the
25 transmission rate at block 814. The amount of streaming data requested to be

1 surged (i.e. transmitted at a higher transmission rate than the initial transmission
2 rate) may be represented in one of several ways. In one implementation, the client
3 604 may request the server 602 to send “*n*” seconds of (future) streaming data at
4 an increased rate. The increased rate may be specifically mentioned (e.g.
5 200Kbps) or the request may denote a transmission rate increase of, for example,
6 ten percent (10%).

7 In another implementation, the client 604 requests that “*x*” number of data
8 packets be transmitted at an increased rate. Again, the increased rate may be
9 specified in a number of ways. In yet another implementation, the client 604 may
10 request the server 602 to send “*y*” bytes of data and an increased rate. Any method
11 by which the client 604 can communicate to the server 602 a request for a
12 particular amount of future streaming data may be utilized in accordance with the
13 systems and methods described herein.

14 When the surged data arrives at the client 604, the bandwidth measurement
15 module 636 assesses the viability of the surged transmission rate by any method
16 known in the art and/or described herein (block 816). This typically entails
17 measuring the streaming bit rate and/or the bandwidth available during the surged
18 transmission.

19 To appropriately identify and measure the surged data, the client 604 must
20 be able to identify and distinguish surged data from pre-surge data and post-surge
21 data.

22 As part of the bandwidth measurement process, the client 604 inspects data
23 packets after a surge request to determine when a surged streaming rate begins and
24 ends. In one implementation, the server 602 is configured to set a flag in a first
25 data packet indicating that the data packet is the first data packet in a series of data

1 packets included in the surged transmission. Similarly, the server 602 is also
2 configured to flag a last data packet to indicate that the data packet is the last data
3 packet in the surged transmission. The client 604 is configured to inspect data
4 packets and identify the flagged data packets.

5 In at least one implementation, the client 604 is configured to request a
6 surge beginning at a specific time for a specific duration. In such an
7 implementation, the client 604 is also configured to identify a time stamp
8 associated with a data packet to determine that the surge has begun, and to identify
9 a time stamp associated with a subsequent data packet to determine that the surge
10 has concluded.

11 If the client 604 has requested a surge of “x” number of data packets, then
12 the client 604 may be configured to identify a flag indicating a data packet is an
13 initial data packet of the surge and to identify the conclusion of the surge by
14 determining when “x” number of data packets have been received by the client.

15 If the client 604 has requested a surge of “y” bytes of data, then the client
16 604 may be configured to identify a flag indicating a data packet is an initial data
17 packet of the surge and to identify the conclusion of the surge by determining
18 when “y” number of bytes have been received.

19 In another implementation, the server 602 may flag each data packet that is
20 surged, i.e. identify the surged data packets as being back to back data packets to
21 be used for bandwidth measurement purposes. In this case, the client 604 is
22 configured to identify each flagged data packet as a surged data packet.

23 In yet another implementation, a separate packet (sometimes known as a
24 “sender report”) is sent that specifies sequence numbers of a first surged packet
25 and a last surged packet. (Each data packet includes a “sequence number” field,

1 which is a number that increments by one for each packet). The surged data
2 packets can be identified from this information.

3 At block 818, the control module 634 determines if the surged rate is viable
4 for continuing for the remaining streaming data. To do this, the control module
5 634 first determines if the measured streaming rate (block 816) is greater than or
6 equal to the current bit rate. If so, the control module 634 then determines if the
7 bit rate of the next higher stream that would be selected if the control module 634
8 requested a stream change is less than or equal to the measured rate (block 816).
9 In other words, if a stream change request would result in a new bit rate to exceed
10 that measured in block 816, then the request should not be made. In other words,
11 in this instance, the surged rate is not viable for continuing for the remaining
12 streaming data ("No" branch, block 818).

13 If the surged rate is not viable for continuing for the remaining streaming
14 data ("No" branch, block 818), then subsequent streaming data is received at the
15 same rate (at block 810) if there is more streaming data available ("Yes" branch,
16 block 824). The process terminates if no further streaming data is available ("No"
17 branch, block 824).

18 If the surged rate is viable for subsequent streaming ("Yes" branch, block
19 818), then the client 604 requests the server 602 to change streams to a next higher
20 bit rate stream in the MBR file (block 820). The streaming bit rate that was
21 measured during the surge (block 816) is then recorded in the history module 638
22 and associated with the server 602 (block 822).

23 In at least one implementation, to reduce the negative impact of a bad
24 measurement taken at block 816, the measured bit rate is combined with the bit
25 rates stored in the history module 638 to derive a new median bit rate. This new

1 median bit rate is the bit rate that is used to determine if the surged rate (i.e. the
2 new median bit rate) is viable for subsequent streaming. However, the bit rate
3 value that is stored in the history module 638 at block 822 is the actual measured
4 bit rate (i.e. before filtering with the bit rates in the history module).

5 **Increasing/Diminishing the Content Buffer**

6 Network congestion can often leave a client buffer smaller than what is
7 required to comfortably shield the user from future network jitter and congestion.
8 To counteract this problem, the client 604 may increase - or grow - the content
9 buffer 648 during a streaming rate surge.

10 During any streaming rate surge the data is received at an increased rate but
11 may not be played at a similar rate. In other words, the content may be executed
12 at a current rate while the surge is in process. The surged data may be used in
13 such instances to grow the content buffer 648 of the client 604 and provide
14 additional protection against artifacts that can detract from the user experience.

15 Likewise, the client 604 may also diminish the size of the content buffer
16 648 in the event that there is more data in the content buffer 648 than is required
17 by the client 604.

18 **Exemplary Computing Environment**

19 Fig. 8 illustrates an example of a suitable computing environment 920 on
20 which the exemplary systems and methods for midstream bandwidth
21 determination may be implemented.

22 Exemplary computing environment 920 is only one example of a suitable
23 computing environment and is not intended to suggest any limitation as to the
24 scope of use or functionality of the exemplary bw-meter. Neither should the
25 computing environment 920 be interpreted as having any dependency or

1 requirement relating to any one or combination of components illustrated in the
2 exemplary computing environment 920.

3 The exemplary midstream bandwidth measurement techniques are
4 operational with numerous other general purpose or special purpose computing
5 system environments or configurations. Examples of well known computing
6 systems, environments, and/or configurations that may be suitable for use with the
7 exemplary systems and methods include, but are not limited to, personal
8 computers, server computers, thin clients, thick clients, hand-held or laptop
9 devices, multiprocessor systems, microprocessor-based systems, set top boxes,
10 programmable consumer electronics, wireless phones, wireless communication
11 devices, network PCs, minicomputers, mainframe computers, distributed
12 computing environments that include any of the above systems or devices, and the
13 like.

14 The exemplary midstream bandwidth measurement systems and methods
15 may be described in the general context of computer-executable instructions, such
16 as program modules, being executed by a computer. Generally, program modules
17 include routines, programs, objects, components, data structures, etc. that perform
18 particular tasks or implement particular abstract data types. The exemplary
19 midstream bandwidth measurement techniques may also be practiced in
20 distributed computing environments where tasks are performed by remote
21 processing devices that are linked through a communications network. In a
22 distributed computing environment, program modules may be located in both local
23 and remote computer storage media including memory storage devices.

24 As shown in Fig. 7, the computing environment 920 includes a general-
25 purpose computing device in the form of a computer 930. The components of

1 computer 920 may include, by are not limited to, one or more processors or
2 processing units 932, a system memory 934, and a bus 936 that couples various
3 system components including the system memory 934 to the processor 932.

4 Bus 936 represents one or more of any of several types of bus structures,
5 including a memory bus or memory controller, a peripheral bus, an accelerated
6 graphics port, and a processor or local bus using any of a variety of bus
7 architectures. By way of example, and not limitation, such architectures include
8 Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA)
9 bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA)
10 local bus, and Peripheral Component Interconnects (PCI) buss also known as
11 Mezzanine bus.

12 Computer 930 typically includes a variety of computer readable media.
13 Such media may be any available media that is accessible by computer 930, and it
14 includes both volatile and non-volatile media, removable and non-removable
15 media.

16 In Fig. 7, the system memory includes computer readable media in the form
17 of volatile, such as random access memory (RAM) 940, and/or non-volatile
18 memory, such as read only memory (ROM) 938. A basic input/output system
19 (BIOS) 942, containing the basic routines that help to transfer information
20 between elements within computer 930, such as during start-up, is stored in ROM
21 938. RAM 940 typically contains data and/or program modules that are
22 immediately accessible to and/or presently be operated on by processor 932.

23 Computer 930 may further include other removable/non-removable,
24 volatile/non-volatile computer storage media. By way of example only, Fig. 7
25 illustrates a hard disk drive 944 for reading from and writing to a non-removable,

1 non-volatile magnetic media (not shown and typically called a “hard drive”), a
2 magnetic disk drive 946 for reading from and writing to a removable, non-volatile
3 magnetic disk 948 (e.g., a “floppy disk”), and an optical disk drive 950 for reading
4 from or writing to a removable, non-volatile optical disk 952 such as a CD-ROM,
5 DVD-ROM or other optical media. The hard disk drive 944, magnetic disk drive
6 946, and optical disk drive 950 are each connected to bus 936 by one or more
7 interfaces 954.

8 The drives and their associated computer-readable media provide
9 nonvolatile storage of computer readable instructions, data structures, program
10 modules, and other data for computer 930. Although the exemplary environment
11 described herein employs a hard disk, a removable magnetic disk 948 and a
12 removable optical disk 952, it should be appreciated by those skilled in the art that
13 other types of computer readable media which can store data that is accessible by a
14 computer, such as magnetic cassettes, flash memory cards, digital video disks,
15 random access memories (RAMs), read only memories (ROM), and the like, may
16 also be used in the exemplary operating environment.

17 A number of program modules may be stored on the hard disk, magnetic
18 disk 948, optical disk 952, ROM 938, or RAM 940, including, by way of example,
19 and not limitation, an operating system 958, one or more application programs
20 960, other program modules 962, and program data 964.

21 A user may enter commands and information into computer 930 through
22 input devices such as keyboard 966 and pointing device 968 (such as a “mouse”).
23 Other input devices (not shown) may include a microphone, joystick, game pad,
24 satellite dish, serial port, scanner, or the like. These and other input devices are
25 connected to the processing unit 932 through an user input interface 970 that is

1 coupled to bus 936, but may be connected by other interface and bus structures,
2 such as a parallel port, game port, or a universal serial bus (USB).

3 A monitor 972 or other type of display device is also connected to bus 936
4 via an interface, such as a video adapter 974. In addition to the monitor, personal
5 computers typically include other peripheral output devices (not shown), such as
6 speakers and printers, which may be connected through output peripheral interface
7 975.

8 Computer 930 may operate in a networked environment using logical
9 connections to one or more remote computers, such as a remote computer 982.
10 Remote computer 982 may include many or all of the elements and features
11 described herein relative to computer 930.

12 Logical connections shown in Fig. 7 are a local area network (LAN) 977
13 and a general wide area network (WAN) 979. Such networking environments are
14 commonplace in offices, enterprise-wide computer networks, intranets, and the
15 Internet.

16 When used in a LAN networking environment, the computer 930 is
17 connected to LAN 977 network interface or adapter 986. When used in a WAN
18 networking environment, the computer typically includes a modem 978 or other
19 means for establishing communications over the WAN 979. The modem 978,
20 which may be internal or external, may be connected to the system bus 936 via the
21 user input interface 970, or other appropriate mechanism.

22 Depicted in Fig. 7, is a specific implementation of a WAN via the Internet.
23 Over the Internet, computer 930 typically includes a modem 978 or other means
24 for establishing communications over the Internet 980. Modem 978, which may
25 be internal or external, is connected to bus 936 via interface 970.

1 In a networked environment, program modules depicted relative to the
2 personal computer 930, or portions thereof, may be stored in a remote memory
3 storage device. By way of example, and not limitation, Fig. 7 illustrates remote
4 application programs 989 as residing on a memory device of remote computer
5 982. It will be appreciated that the network connections shown and described are
6 exemplary and other means of establishing a communications link between the
7 computers may be used.

8 **Exemplary Operating Environment**

9 Fig. 8 illustrates an example of a suitable operating environment 920 in
10 which the midstream bandwidth determination techniques may be implemented.
11 Specifically, the midstream bandwidth determination is implemented by any
12 program 960-962 or operating system 958 in Fig. 8.

13 The operating environment is only an example of a suitable operating
14 environment and is not intended to suggest any limitation as to the scope of use of
15 functionality of the midstream bandwidth measurement techniques described
16 herein. Other well known computing systems, environments, and/or
17 configurations that may be suitable for use with the described systems and
18 methods include, but are not limited to, personal computers, server computers,
19 hand-held or laptop devices, multiprocessor systems, microprocessor-based
20 systems, programmable consumer electronics, network PCs, minicomputers,
21 mainframe computers, distributed computing environments that include any of the
22 above systems or devices, and the like.

23 **Computer-Executable Instructions**

24 An implementation of the exemplary midstream bandwidth measurement
25 systems and methods may be described in the general context of computer-

1 executable instructions, such as program modules, executed by one or more
2 computers or other devices. Generally, program modules include routines,
3 programs, objects, components, data structures, etc. that perform particular tasks
4 or implement particular abstract data types. Typically, the functionality of the
5 program modules may be combined or distributed as desired in various
6 embodiments.

7 **Computer Readable Media**

8 An implementation of the exemplary midstream bandwidth measurement
9 techniques may be stored on or transmitted across some form of computer
10 readable media. Computer readable media can be any available media that can be
11 accessed by a computer. By way of example, and not limitation, computer
12 readable media may comprise computer storage media and communications
13 media.

14 Computer storage media include volatile and non-volatile, removable and
15 non-removable media implemented in any method or technology for storage of
16 information such as computer readable instructions, data structures, program
17 modules, or other data. Computer storage media includes, but is not limited to,
18 RAM, ROM, EEPROM, flash memory or other memory technology, CD-ROM,
19 digital versatile disks (DVD) or other optical storage, magnetic cassettes, magnetic
20 tape, magnetic disk storage or other magnetic storage devices, or any other
21 medium which can be used to store the desired information and which can be
22 accessed by a computer.

23 Communication media typically embodies computer readable instructions,
24 data structures, program modules, or other data in a modulated data signal such as
25 carrier wave or other transport mechanism and included any information delivery

1 media. The term “modulated data signal” means a signal that has one or more of
2 its characteristics set or changed in such a manner as to encode information in the
3 signal. By way of example, and not limitation, communication media includes
4 wired media such as a wired network or direct-wired connection, and wireless
5 media such as acoustic, RF, infrared, and other wireless media. Combinations of
6 any of the above are also included within the scope of computer readable media.
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25

1 **Conclusion**

2 Although the subject matter has been described in language specific to
3 structural features and/or methods, it is to be understood that the invention defined
4 by the appended claims is not necessarily limited to the specific features or
5 methods described herein. Rather, the specific features and methods are disclosed
6 as exemplary forms of implementing the claimed systems and methods.

7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25